The Prophet Inequality

CS 1951k/2951z

2020-03-11

We describe the prophet inequality, and several simple near-optimal mechanisms that follow as an immediate consequence.

1 *The Prophet Inequality*

The prophet inequality[1] is a surprising result that lower bounds the expected reward[2] one can obtain in the following game: You are the leader of an investment team. Your team is in possession of some assets that your most trustworthy monetary advisor claims will be worthless after n days. Panicked, you decide to liquidate these assets within the next n days, but want to do so while still maximizing your rewards from selling the assets. Fortunately, your quant team comprises CSCI 1951k/2951z alumni who are accurately able to back out the *exact* distributions F_i of the assets' market values π_i on day i, for all $1 \leq i \leq n$.[3] Equipped with this information, it is your job to make the executive decision as to when you will sell the assets.

Once again:

1. The market value π_i on day i is drawn from some distribution F_i.
2. At the start of day i, you observe the market value π_i.
3. You then have two choices: cash out now, or wait.
4. If you cash out now, you walk away with a reward of π_i.
5. If you wait, then you can no longer sell your assets at price π_i, and you will face the same decision the next day.
6. After n days, your assets become worthless.
7. Your task is to devise a strategy for deciding whether you should sell on day i or wait another day, after observing π_i.

This problem is an instance of an **optimal stopping problem**, where the goal is to devise a **stopping rule** that tells you when to *stop*, which in our setup means when to cash out. Another famous stopping problem is the **secretary problem**. There, you are interviewing secretaries,[4] and your goal is to decide when to stop interviewing so as to maximize the probability of hiring the best secretary, without assuming the order in which applicants arrive is known *a priori*, or even without assuming any distributional knowledge at all.

[1] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. In J. Kuelbs, editor, *Probability on Banach Spaces*, volume 4, pages 197–266. Dekker, 1978; and Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 11 1984

[2] or utility, or payoff, or etc.

[3] Perhaps surprisingly, market values from one day to the next in this game are independent of one another.

[4] or applicants for any job, really, including spouse

Optimal stopping rules can be computed via dynamic programming when, n is small. For certain variants of the problem, these rules take the form: reject the first m applicants, and then accept the first applicant thereafter who is preferable to the first m.

In this lecture, we will not derive an optimal stopping rule, but rather an approximately-optimal one. Our rule will take the form of a thresholding strategy, where we accept the first applicant whose value is above a pre-computed threshold. Note that computing this threshold depends on knowledge of the distributions F_i of the random variables. Implicit in the F_i notation, it should be clear that we are also assuming independence.

Imagine a clairvoyant *prophet* who can see all of the π_i's in advance. With this information in hand, the optimal strategy is straightforward: choose the largest $\pi_i \geq 0$. In what follows, we will be comparing the expected reward of a very simple thresholding strategy to that of the prophet. The simple thresholding strategy is: fix a value t, and then sell the assets on the first day, if any, that $\pi_i \geq t$.[5]

[5] Note that this strategy is clearly sub-optimal, as you would never retain the assets on the very last day; nevertheless, our simple strategy never sells the assets if π_i never exceeds t.

We claim that for an appropriate choice of t, this simple thresholding strategy is a 2-approximation of the optimal: i.e., the prophet's strategy. In fact, there are two distinct choices of t, both of which achieve this goal. We derive one of them presently—the one which lends itself to the development of approximately-optimal auctions.

2 *Proof of the Prophet Inequality*

Let APX and OPT denote the expected revenue of the thresholding strategy and the prophet, respectively. We will proceed by first lower bounding APX, and then upper bounding OPT.

APX, the expected reward of the threshold strategy, is the sum of the expected values that the assets are sold on each day i, which occurs only only if $\pi_i \geq t$ and $\pi_j < t$, $\forall j < i$: i.e.,

$$\begin{aligned}
\text{APX} &= \sum_{i=1}^{n} \mathbb{E}[\pi_i \mid \pi_i \geq t; \pi_j < t, \forall j < i] \Pr(\pi_i \geq t; \pi_j < t, \forall j < i) \\
&= \sum_{i=1}^{n} \mathbb{E}[\pi_i \mid \pi_i \geq t; \pi_j < t, \forall j < i] \Pr(\pi_i \geq t \mid \pi_j < t, \forall j < i) \Pr(\pi_j < t, \forall j < i) \\
&= \sum_{i=1}^{n} \mathbb{E}[\pi_i \mid \pi_i \geq t] \Pr(\pi_i \geq t) \Pr(\pi_j < t, \forall j < i) \text{ by independence} \\
&= \sum_{i=1}^{n} (\mathbb{E}[\pi_i - t \mid \pi_i \geq t] + t) \Pr(\pi_i \geq t) \Pr(\pi_j < t, \forall j < i) \\
&= \sum_{i=1}^{n} \mathbb{E}[\pi_i - t \mid \pi_i \geq t] \Pr(\pi_i \geq t) \Pr(\pi_j < t, \forall j < i) + t \left(\sum_{i=1}^{n} \Pr(\pi_i \geq t) \Pr(\pi_j < t, \forall j < i) \right)
\end{aligned}$$

Now let p_t denote the probability that, at the threshold t, the assets are not sold after all n days: i.e., $p_t = \Pr(\pi_i < t, \forall 1 \le i \le n)$. Using this notation, the probability that the assets are sold on one of days 1 through n, $\sum_{i=1}^{n} \Pr(\pi_i \ge t) \Pr(\pi_j < t, \forall j < i)$, is $1 - p_t$. Thus, letting $(x)^+ = \max(0, x)$, we obtain a lower bound on APX as follows:

$$\begin{aligned}
\text{APX} &= \sum_{i=1}^{n} \mathbb{E}[(\pi_i - t)^+] \Pr(\pi_j < t, \forall j < i) + t(1 - p_t) \\
&\ge \sum_{i=1}^{n} \mathbb{E}[(\pi_i - t)^+] \Pr(\pi_j < t, \forall 1 \le j \le n) + t(1 - p_t) \\
&= \sum_{i=1}^{n} \mathbb{E}[(\pi_i - t)^+] p_t + t(1 - p_t)
\end{aligned}$$

Having established a lower bound on APX, we proceed to derive an upper bound on OPT. Since the prophet will sell the assets when their price is maximal, the value of the prophet's strategy can be upper bounded as follows:

$$\begin{aligned}
\text{OPT} &= \mathbb{E}\left[\max_i \pi_i\right] \\
&\le \mathbb{E}\left[\max_i \pi_i^+\right] \\
&= \mathbb{E}\left[\max_i (\pi_i^+ - t)\right] + t \\
&\le \mathbb{E}\left[\max_i (\pi_i - t)^+\right] + t \\
&\le \sum_{i=1}^{n} \mathbb{E}\left[(\pi_i - t)^+\right] + t
\end{aligned}$$

The only non-trivial step in this derivation is the second to last, which follows as $\max(\pi_i, 0) - t \le \max(\pi_i, t) - t = \max(\pi_i - t, 0)$.

Note that the above inequality holds for all thresholds t. We choose t such that $p_t = 1/2$ to obtain the desired 2-approximation:

$$\begin{aligned}
\text{APX} &\ge \sum_{i=1}^{n} \mathbb{E}\left[(\pi_i - t)^+\right] p_t + t(1 - p_t) \\
&= 1/2 \left(\sum_{i=1}^{n} \mathbb{E}\left[(\pi_i - t)^+\right] + t \right) \\
&\ge 1/2\, \text{OPT}
\end{aligned}$$

It turns out that another threshold value works just as well as the choice we just derived. This value is $t = 1/2\, \mathbb{E}[\pi^*]$, where $\pi^* = \max_i \pi_i$. A particularly elegant proof of this fact appears in Correa *et al.*,[6] as their proof implies both our choice and this latter choice at once. In particular, they prove $\text{APX} \ge p_t t + (1 - p_t)(\mathbb{E}[\pi^*] - t)$. So, both choosing t either such that $p_t = 1/2$ (as above) and setting $t = 1/2\, \mathbb{E}[\pi^*]$ yields the desired inequality.

[6] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Recent developments in prophet inequalities. *SIGecom Exchanges*, 17(1):61–70, 2018

3 *Applications of the Prophet Inequality*

The prophet inequality can be used to design a 2-approximation of Myerson's optimal auction. This auction works as follows: Given bidders $1, \ldots, n$, with values drawn from *regular* distributions $F_1, \ldots, F_n$, set the reserve price for bidder i to be $\varphi_i^{-1}(t)$, where t is determined by the prophet inequality, and φ_i is bidder i's virtual value function. Then pick two bidders who meet their reserve and allocate the good to the one of the two who bids higher, at a price which is the maximum of their reserve and the other's bid.[7]

[7] Of course, it would be sensible for an auctioneer who seeks to maximize revenue to choose the two highest bidders, but the bound holds for any two bidders who meets their reserve.

Viewed from the point of view of the prophet inequality, bidders turn up (sequentially or simultaneously; their arrival order does not impact the bound), and each bidder i places bid v_i. (We assume incentive compatibility, as we enforce this constraint via Myerson's payments.) Any bidder i who meets their reserve bids such that $v_i \geq \varphi_i^{-1}(t)$; equivalently, $\varphi_i(v_i) \geq t$. Therefore, the revenue APX accrued by an auctioneer that employs a bidder-dependent threshold strategy with threshold $\varphi_i^{-1}(t)$ is lower bounded by $1/2\,\mathrm{OPT}$, where OPT is the revenue of Myerson's optimal auction:

$$
\begin{aligned}
\mathrm{APX} &\geq \sum_{i=1}^{n} \mathbb{E}\left[(\varphi_i(v_i) - t)^+\right] p_t + t(1 - p_t) \\
&= 1/2 \left(\sum_{i=1}^{n} \mathbb{E}\left[(\varphi_i(v_i) - t)^+\right] + t \right) \\
&\geq 1/2 \left(\mathbb{E}\left[\max_i (\phi_i(v_i))^+ \right] \right) \\
&= 1/2 \left(\mathbb{E}\left[\max_i \phi_i(v_i) x_i(\boldsymbol{v}) \right] \right) \\
&= 1/2\,\mathrm{OPT}
\end{aligned}
$$

This auction, while sub-optimal, is simpler than Myerson's optimal auction. It relies on the virtual value function only to set the bidders' reserve prices, and even then, it uses the function only at one specific point, namely $\varphi_i^{-1}(t)$. Moreover, it is more natural, as it can (and probably should; see Sidenote 7) allocate to the highest bidder.

There is also a straightforward connection between the prophet inequality and posted-price mechanisms. Imagine a (matchbox) car salesman who has posted a price on his prized Maserati. Each day, another potential buyer enters his store, admires the car, and buys it or not, depending on their value for the car. They buy it precisely when their value exceeds (or matches) the price.

Assuming the salesman plans to entertain exactly n buyers, one per day, and has knowledge of their value distributions, namely $F_1, \ldots, F_n$, he can post $\varphi_i^{-1}(t)$ as the price for buyer i, where as above, t is determined by the prophet inequality and φ_i is buyer i's virtual value function. By the argument above, this posted price will guarantee him half the revenue of Myerson's optimal auction, assuming the buyers walk away with the car if their value exceeds their price.

As the salesman operates in a sequential setting, perhaps he can earn even more revenue by dynamically updating his threshold strategy. You will explore this question in your homework.

References

[1] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Recent developments in prophet inequalities. *SIGecom Exchanges*, 17(1):61–70, 2018.

[2] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. In J. Kuelbs, editor, *Probability on Banach Spaces*, volume 4, pages 197–266. Dekker, 1978.

[3] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 11 1984.